

VISIT...

LANZAROTE
Caliente.COM

Real-Time Movement Prediction for Improved Control of Neuroprosthetic Devices

Andreas A. C. Thomik¹, David Haber² and A. Aldo Faisal^{1,2,3}, *Member IEEE*

Abstract—Replacing lost hands with prosthetic devices that offer the same functionality as natural limbs is an open challenge, as current technology is often limited to basic grasps by the low information readout. In this work, we develop a probabilistic inference-based method that allows for improved control of neuroprosthetic devices. We observe the behaviour of the undamaged limb to predict the most likely actions of lost limbs. Offline, our algorithm learns movement primitives (e.g. various types of grasps) from a database of recordings from healthy subjects performing everyday activities. Online, it performs Bayesian inference to determine the currently active movement primitive from the observed limbs and estimates the most likely movement of the missing limbs from the training data. We can demonstrate on test data that this two-stage approach yields statistically significantly higher prediction accuracy than linear regression approaches that reconstruct limb movements from their overall correlation structure.

I. INTRODUCTION

Neuroprosthetics offer the chance for people with a lost limb to regain naturalistic mobility and dexterity by controlling a prosthetic replacement (such as an artificial hand) in the same way they would control their own body. Whilst spectacular advances have been made by using direct cortical control from both invasive [1], [2], [3] and non-invasive electrodes [4], [5] as well as using other information sources such as EMG [6], [7], [8], the amount of information we can extract with current technology is still very low and not suited for control of prosthetic devices with the same versatility as the natural limb [9], [10]. This restricts the potential use of neuroprosthetics in everyday activities considerably and often leads to frustration and rejection of the prosthesis by the user [11], [12].

Here, we suggest to enhance the control of neuroprosthetic devices by exploiting two aspects of natural movement: (1) modularity, by which complex movements are a combination of a finite set of simpler movements [13], [14]; and (2) stereotypical patterns of correlation in the variability structure of joints movement during everyday tasks [15], [16].

We base our method on a large data set of complex hand movements which represents more than 6 hours worth of data from different subjects performing a variety of tasks [17]. We extract and cluster movement primitives from this data by apply a segmentation algorithm we developed. These allow us to accurately complement partial observations (e.g.

movement of three fingers) and predict the behaviour of the unobserved joints (e.g. remaining two fingers) in real-time by taking into account the probabilistic structures of the most representative clusters. We show that by considering movement primitives (or segments) instead of the entire, un-segmented time-series, we can significantly improve the prediction accuracy for the movement of non-observed joints with significantly less training data.

II. METHODS

A. Data

We use hand movement data collected from 7 healthy subjects performing everyday tasks such as opening doors, using the phone, eating, etc. [17] The data was collected using a CyberGlove I (CyberGlove Systems, San Diego, CA) which samples movements at 80 Hz with a resolution of 8 bits per sensor. The device outputs 18-dimensional time-series representing the bending of all joints of the fingers with exception of the distal interphalangeal joint (DIP), as well as flexion/extension and adduction/abduction of the wrist. Before further analysis, the raw data is smoothed using a second-order Savitzky-Golay filter with a running window of 5 data points to remove discontinuities induced by the A/D converter. Finally, sensor values relating to wrist position are discarded to obtain a 16-dimensional data set of pure finger movements. The subsequent analysis is performed on the velocity values of the data.

B. Segmentation

We segment the data from all subjects and tasks using the method described in (D Haber et al. - *in review*). The algorithm first performs PCA to reduce the data dimensionality before segmenting the data using an efficient Bayesian changepoint detection algorithm and a Bayesian prior over movement duration. Based on the changepoints that were computed we extract movement segments from the original data. The segmentation algorithm scales linearly with the number of data points and is near constant in dimensionality of the data which allows for the processing of very large data sets as collected from motion capture equipment over a prolonged period of time.

C. Clustering

A sequential process is used to measure similarity between two segments. The segments are resampled to a default movement duration to ensure consistency in the number of data points per segment. We then normalise all segments to have amplitudes in the interval $[-1, 1]$ by dividing each

¹Department of Bioengineering, Imperial College London, SW7 2AZ London, UK andreas.thomik10 at imperial.ac.uk

²Brain & Behaviour Lab, Dept. of Computing, Imperial College London, SW7 2AZ London, UK david.haber09 at imperial.ac.uk

³MRC Clinical Sciences Centre, Hammersmith Hospital Campus, W12 0NN London, UK a.faisal at imperial.ac.uk




Fig. 1. Conceptual design of the prosthesis controller. **Left hand side:** A limited amount of information is observed (here, two fingers) and inference is performed to evaluate which movement primitive most likely generated the observed data. **Right hand side:** The parameters of the most likely movement primitive are used to estimate the missing dimensions and the prosthetic replacement is actuated accordingly. The numbers on the fingers indicate their ID.

segment's amplitude by its absolute maximum. Finally, the similarity between clusters is obtained by computing the maximum cross-correlation (MXC) of all possible pairs of segments over a limited time-lag corresponding to 20% of the segment length. This corrects for slight misalignments of movement within the segments. We found that alternative methods such as the Pearson correlation coefficient were less accurate indicators of similarity as they are highly affected by even marginal shifts in time between segments that represent the same type of movement.

Clustering of segments is performed using an iterative method which, starting with an arbitrary segment, finds all other segments which are more similar than a given threshold t and repeats this with the found segments until the cluster converges to a stable state. Visual inspection of clustered elements showed clear resemblance in the movement executed. A slight drawback of this method is that it generates a number of clusters with a single segment as they are deemed too different from any other segment and therefore seen as truly unique by the algorithm. The number of these unique clusters can be regulated by t . These unique segments can be caused by either (1) unnatural manipulation of objects or (2) erroneous segmentation of the time-series by the algorithm in the pre-processing step. These segments were ignored in further analysis.

D. Inference

To infer which cluster a given data point may belong to, we fit all segment clusters with a mixture of Gaussians in order to obtain a probabilistic representation of the data. We choose the optimal number of components for each mixture of Gaussians using the Bayesian Information Criterion. Consequently, the likelihood of the data at time t being generated by a particular cluster S_i is given by

$$\pi_i = p(S_i|x_{1:t}^{(obs)}) = \frac{p(x_{1:t}^{(obs)}|S_i) \times p(S_i)}{p(x_{1:t}^{(obs)})} \quad (1)$$

where $p(S_i)$ can be estimated directly from the training data and $x_{1:t}$ is the data observed so far. We use a weighted sum of probabilities for the data points observed in the

past such that the influence of past data points decays exponentially over time.

E. Prediction

To predict the unobserved dimensions, we fit a regression with a linear combination of 50 Gaussian basis functions between the observed and unobserved dimensions for each cluster in the data set offline. Upon observing partial movement behaviour, we can compute the unobserved dimensions online and in real-time. We select the most likely cluster from the previously computed π_i distribution in a winner-takes-all fashion and estimate the hidden trajectory on the basis of the regression parameters fitted to that particular cluster. The time required to predict the unobserved movement behaviour at a given time-step is therefore only dependent on the number of clusters considered during the prediction step.

III. RESULTS

The following section presents the results of the segmentation, clustering and prediction steps of our method. We show that by exploiting the unique characteristics of each movement primitive cluster we are capable of predicting the unobserved data much more precisely than by taking into consideration the entirety of the movement data.

1) *Movement Primitives:* The segmentation of the data into movement primitives by the algorithm produced a total of 3710 primitives with an average movement duration of 3.28 seconds. Note that when a block of data is segmented, the first and last segments are discarded as they are potentially incomplete. Clustering of the movement primitives as described in the Methods section produced 2279 clusters when the similarity threshold was set to $t = 0.8$. We note that of these clusters, a vast majority (88.85%) were singleton clusters, i.e. segments for which the similarity with any other segment was less than t .

To simplify the procedure of estimating the regression parameters and to limit the influence of the singleton segments, we only consider movement primitive clusters which account for at least 0.2% of the total number of data points collected. This means that out of the 2279 clusters that are initially




Fig. 2. (A) Raw data of a subject plugging and unplugging a electric plug 10 times. Vertical black dashed lines indicate segmentation points as determined by our algorithm. (B) Similarity between the segments from (A) and the general movement primitives extracted from a larger data set. The brighter the case, the more similar the segments are, with black no similarity, and white being identity. (C) The sequence of the most likely movement primitive active at each given time segment. See text for details.

present, only 34 (1.5% of the clusters, 25.6% of the all data points) were used to predict the unobserved data. In Fig. 2 we show the similarity of the cluster core segments with the segments obtained from a longer recording in which the subject picked up and put down a bag 10 times. The raw data and segmentation are shown in Fig. 2A. Fig. 2B shows the similarity of the cluster centroids with the extracted segments as given by maximum cross-correlation. Note that the data shown here was not used to form the clusters against which it is being compared. Fig. 2C presents the segments with the highest similarity. The repetitive nature of the task becomes evident from the alternating order of underlying primitives.

2) *Prediction*: Estimation of the unobserved data was a two step process in which (1) the movement primitive cluster which has most likely generated the observed data is selected and (2) unobserved dimensions are estimated using a regression particularly fitted to the data of that movement primitive cluster (see Methods for details). We note that when the most likely generating movement primitive changed between time-step t and $t+1$, we sometimes observe a considerable change in the estimate of x_t^{hid} and x_{t+1}^{hid} . This can be explained by the statistical nature of primitive inference which sometimes leads to the wrong primitive being used for prediction. We safeguarded ourselves from such crude mistakes by fitting an exponential distribution to the absolute change in value between two time-steps of the unobserved data in our training set. If a jump of magnitude $|\hat{x}_t - \hat{x}_{t+1}|$ or larger had a likelihood less than 0.10%, we assumed that the inference had failed and re-iterated the movement primitive inference

the current estimate discarded. This successfully prevented jumps in the estimate of unobserved dimensions whilst additionally refining the inference step for the underlying movement primitive.

Using this method, we tested two different scenarios (see Fig. 3): (1) only fingers 1 and 2 (for finger identity, see Fig. 1) were observed and the movement of fingers 3-5 estimated (Fig. 3A); and (2) thumb, index and middle finger were observed and the movement of ring and little finger reconstructed (Fig. 3B). Over all data sets at our disposal, we thus obtained an average $R^2 = 0.19 (\pm 0.031)$ for the first and $R^2 = 0.25 (\pm 0.023)$ for the latter case (see Fig. 3, white columns). Interestingly, the Pearson correlation coefficient (ρ) between the ground truth and the reconstructed data on the same data set was significantly higher than the R^2 values (Fig. 3A: $\rho = 0.47 (\pm 0.013)$; 3B: $\rho = 0.49 (\pm 0.014)$). The discrepancy between R^2 and ρ suggests that although we correctly identify the shape of the unobserved time-series, the scaling is not perfect. This can be understood by considering that movements can be performed at varying speeds and the regression only takes into account the average movement velocity.

In a second test, we compared the results from our method with the simpler approach of estimating the movement of the missing fingers by calculating the regression parameters over the unsegmented data from 5 of our 7 subjects and estimating the missing dimensions on the remaining two subjects (Fig. 3, gray columns). We find that although our novel method only takes into account a fraction of the data available, it

performs significantly better (Student t-test, $p < 0.05$) than the second approach which takes into account all the data – this becomes particularly evident when trying to estimate the movement of three fingers.




Fig. 3. Performance of our method (white) against regression over all the observed data (gray). (A) Estimate of the trajectory of fingers 3–5 from the observation of fingers 1 and 2 (B) Estimate of the trajectory of fingers 4 and 5 from the observation of fingers 1–3. Bars indicate standard error. *: $p < 0.05$. See text for details.

IV. CONCLUSIONS

In this paper we presented a method which is capable of accurately predicting the movement of missing finger joints from the remaining fingers. Instead of relying on direct information such as EMG of the muscles, we decompose training data into movement primitives and exploit the typical correlation patterns which are characteristic for those primitives. A probabilistic representation of a large data set containing everyday movement segments is generated during a pre-processing step and allows for movement prediction in real-time. We successfully tested our method by reconstructing the movement of missing fingers on a large data set of everyday tasks [17] and showed significantly better performance than the un-segmented approach.

We propose to integrate this method into the conceptual design of improved neuroprosthetic devices which on top of decoding action intention from a variety of brain-machine interfaces (BMIs) such as EMG, MMG, intra-cortical electrodes or 3D gaze-tracking can make use of prior information about human movement statistics to improve the control [18], [19], [20]. This would be of particular interest if we consider that it may be possible to predict the movement of the fingers by only observing the movement of shoulder and elbow or the other hand, such as would be the case in a hand amputee. A prior over likely trajectories is of particular interest as it allows to make more efficient use of the limited information obtained with current brain readout mechanisms.

Based on the promising results laid out in this study, we suggest to improve the numerical tools for segmentation and regression and extend this approach to all limbs.

V. ACKNOWLEDGEMENTS

We thank Jovana Belic for collecting the data set. AACT is supported by the National Research Fund, Luxembourg (1229297). AAF is supported by an HFSP Program Grant (RPG00022/2012) and the Human Frontiers in Science Program Grant (HFSP RPG00022/2012).

REFERENCES

- [1] M Velliste et al. Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198):1098–1101, 2008.
- [2] LR Hochberg et al. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442(7099):164–171, 2006.
- [3] LR Hochberg et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375, 2012.
- [4] JR Wolpaw and DJ McFarland. Multichannel eeg-based brain-computer communication. *Electroencephalogr Clin Neurophysiol*, 90(6):444–449, 1994.
- [5] JR Wolpaw and DJ McFarland. Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proc Natl Acad Sci USA*, 101(51):17849–17854, 2004.
- [6] S Bitzer and P van der Smagt. Learning emg control of a robotic hand: Towards active prostheses. In *Proc. IEEE ICRA*, pages 2819–2823. IEEE, 2006.
- [7] TA Kuiken et al. Targeted reinnervation for enhanced prosthetic arm function in a woman with a proximal amputation: a case study. *Lancet*, 369(9559):371–380, 2007.
- [8] TA Kuiken et al. Targeted muscle reinnervation for real-time myoelectric control of multifunction artificial arms. *J Am Med Assoc*, 301(6):619–628, 2009.
- [9] WW Abbott and AA Faisal. Ultra-low-cost 3d gaze estimation: an intuitive high information throughput compliment to direct brain-machine interfaces. *J Neural Eng*, 9(4):046016, 2012.
- [10] WW Abbott et al. Large-field study of ultra low-cost, non-invasive task level bmi. In *Neural Engineering (NER), 2013 6th International IEEE/EMBS Conference on*, 2013.
- [11] RA Roeschlein and E Domholdt. Factors related to successful upper extremity prosthetic use. *Prosthet Orthot Int*, 13(1):14–18, 1989.
- [12] J Davidson. A survey of the satisfaction of upper limb amputees with their prostheses, their lifestyles, and their abilities. *J Hand Ther*, 15(1):62–70, 2002.
- [13] E Bizzi et al. Modular organization of motor behavior in the frog's spinal cord. *Trends Neurosci*, 18(10):442–446, 1995.
- [14] KA Thoroughman and R Shadmehr. Learning of action through adaptive combination of motor primitives. *Nature*, 407(6805):742–747, 2000.
- [15] AAC Thomik and AA Faisal. Deriving motion primitives from naturalistic hand movements for neuroprosthetic control. *BCCN 2012 Munich*, page 201, 2012.
- [16] A Aldo Faisal et al. Noise in the nervous system. *Nat Rev Neurosci*, 9(4):292–303, 2008.
- [17] J Belic and AA Faisal. The structured variability of finger motor coordination in daily tasks. *BMC Neurosci*, 2011.
- [18] WW Abbott and AA Faisal. Ultra-low-cost 3d gaze estimation: an intuitive high information throughput compliment to direct brain-machine interfaces. *Journal of neural engineering*, 9(4):046016, 2012.
- [19] C Gavriel and AA Faisal. Wireless kinematic body sensor network for low-cost neurotechnology applications "in-the-wild". In *Neural Engineering (NER), 2013 6th International IEEE/EMBS Conference on*, 2013.
- [20] S Fara et al. Robust, ultra low-cost mmg system with brain-machine-interface applications. In *Neural Engineering (NER), 2013 6th International IEEE/EMBS Conference on*, 2013.